

Technical Report for ICCV 2023 Visual Continual Learning Challenge: Continuous Test-time Adaptation for Semantic Segmentation

Damian Sójka^{1,2}, Yuyang Liu^{3,4}, Dipam Goswami^{5,6}, Sebastian Cygert^{2,7}, Bartłomiej Twardowski^{2,5,6},
and Joost van de Weijer^{5,6}

¹Poznań University of Technology, ²IDEAS NCBR, ³Shenyang Institute of Automation,
⁴Chinese Academy of Sciences, ⁵Computer Vision Center, ⁶Autonomous University of Barcelona,
⁷Gdańsk University of Technology

1. Introduction

The goal of the challenge is to develop a test-time adaptation (TTA) method, which could adapt the model to gradually changing domains in video sequences for semantic segmentation task. It is based on a synthetic driving video dataset - SHIFT [5]. The source model is trained on images taken during daytime in clear weather. Domain changes at test-time are mainly caused by varying weather conditions and times of day. The TTA methods are evaluated in each image sequence (video) separately, meaning the model is reset to the source model state before the next sequence. Images come one by one and a prediction has to be made at the arrival of each frame. Each sequence is composed of 401 images and starts with the source domain, then gradually drifts to a different one (changing weather or time of day) until the middle of the sequence. In the second half of the sequence, the domain gradually shifts back to the source one. Ground truth data is available only for the validation split of the SHIFT dataset, in which there are only six sequences that start and end with the source domain. We conduct an analysis specifically on those sequences. Ground truth data for test split, on which the developed TTA methods are evaluated for leader board ranking, are not publicly available.

The proposed solution secured a 3rd place in a challenge and received an innovation award. Contrary to the solutions that scored better, we did not use any external pretrained models or specialized data augmentations, to keep the solutions as general as possible. We have focused on analyzing the distributional shift and developing a method that could adapt to changing data dynamics and generalize across different scenarios.

2. Problem Analysis

We conduct an extensive analysis using the semantic segmentation source model DeepLabv3+ [1] with ResNet50 [2] backbone. We utilize model weights provided by the challenge organizers.

Firstly, we check how the domain shift influences unchanged source model performance over the span of the sequence. Figure 1 shows the mean intersection over union (mIoU) for each image in six sequences from the validation split. As expected, the performance degrades until the middle of the sequence, where the domain shift is the most drastic, and starts increasing toward the end of the sequence, as the domain comes back to the source one. However, the change in mIoU is gradual for sequences in which the domain change is in the form of weather conditions (clear to rainy or foggy), and more abrupt for videos with domain change to night. It might be caused by drastic changes in lighting conditions during the night. This suggests that the developed method should be flexible and be able to adapt the model to both gradual and abrupt changes.

To further analyze the nature of domain shift, we decided to inspect the shift in data distribution $\phi = (\mu, \sigma)$, where μ is a mean of data and σ represents a standard deviation. We examine the distance between the distributions of source data ϕ^S and test images ϕ_t^T for each frame at time t . We utilize symmetric Kullback-Leibler divergence as a distance metric $D(\phi^S, \phi_t^T)$:

$$D(\phi^S, \phi_t^T) = \frac{1}{C} \sum_{i=1}^C KL(\phi_i^S || \phi_{t,i}^T) + KL(\phi_{t,i}^T || \phi_i^S) \quad (1)$$

where C is the number of channels.

The distance plot is presented in Figure 2. It can be seen that the significance of the distribution shift varies between

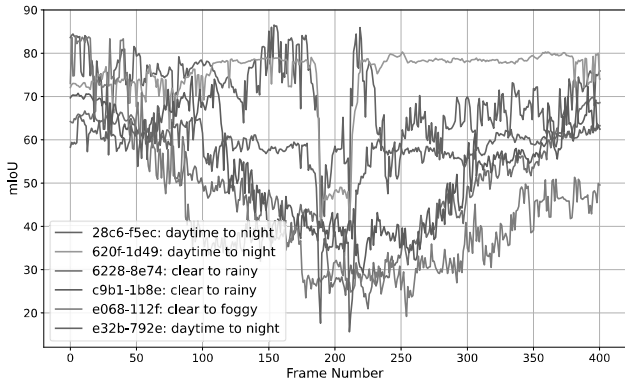


Figure 1. The plot of $mIoU$ between ground truth and source model predictions for each frame in six different sequences from the validation split.

sequences. Changes in time of day to night time cause the test image distribution to greatly drift from the source data distribution. On the other hand, differences in weather conditions do not influence the distribution changes significantly. It shows that the TTA method should be able to handle different distributions of data and adjust the normalization process of the model accordingly, especially while using batch normalization (BN).

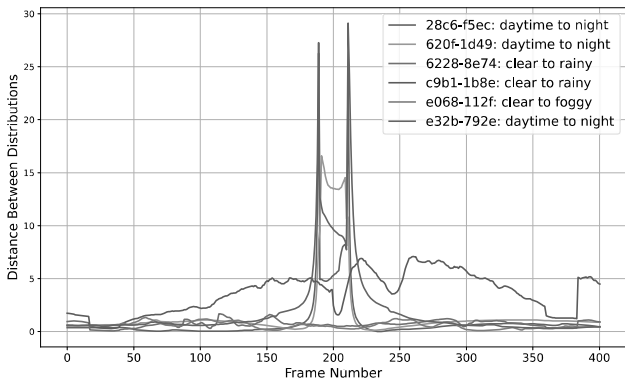


Figure 2. The plot of symmetric Kullback-Leibler divergence as a distance metric between source training data distribution and the distribution of each frame in six different sequences from the validation split.

Lastly, Figure 3 shows the mean entropy of source model predictions for each frame of the sequences. Predictions' entropy increases with increasing domain shift. Moreover, the trend is highly similar to the $mIoU$ plot in Figure 1. It suggests that entropy might be a relatively useful metric for evaluating the model performance and the degree of domain shift during test-time.

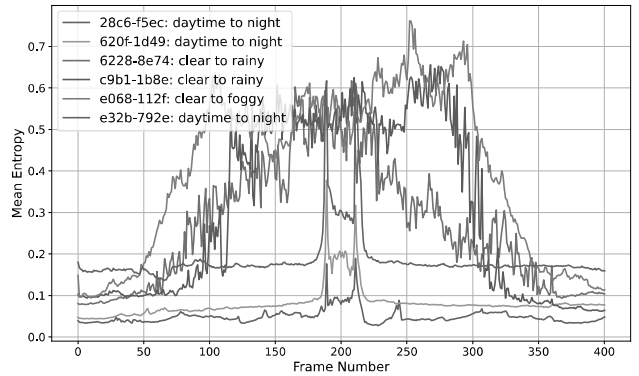


Figure 3. The plot of mean entropy of source model predictions for each frame in six different sequences from the validation split.

3. Baselines

There are two baseline TTA methods implemented by the challenge organizers: TENT [7] and CoTTA [8]. They use two different adaptation methods. TENT uses prediction entropy minimization to update only batch normalization weights. CoTTA is based on adapting the student model with pseudo-labels generated by the teacher model. The teacher is updated by the exponential moving average of student's weights. To prevent performance degradation, it additionally uses stochastic model restoration, where randomly selected weights are reset to the source model state. Table 1 shows their performance using local evaluation on six sequences from the validation split. DeepLabv3+ [1] model with weights provided by the organizers is used. TENT outperformed the CoTTA method in the challenge setting. Therefore, we choose entropy minimization used in TENT as our base adaptation method and build upon it.

Table 1. $mIoU$ and overall metrics of fixed source model, TENT [7] and CoTTA [8] baselines. The results are from the local evaluation on the validation split. The learning rate is equal to 0.00006/8.

Method	$mIoU$	overall
Source model	54.4	12.9
TENT [7]	58.1	41.0
CoTTA [8]	56.2	12.0

4. Our Method

Our method is composed of the base adaptation method from TENT [7] - entropy minimization, and three additional modifications. Firstly, considering the experimental results from Figure 2, we utilize the dynamic BN statistics update method, described in Section 4.1. Moreover, to make the entropy minimization process more reliable, we filter the uncertain pixels from the minimization process by the value of the entropy of their prediction. We depict this approach

in Section 4.2. Lastly, the original TENT method for classification tasks adapted only BN weights of the whole model. For segmentation, we only adapt the BN weights of the backbone (ResNet50), leaving the segmentation head fixed. We show the advantage of this approach experimentally in Section 5.4.

4.1. Dynamic Batch Normalization Statistics Update

Due to domain shift, state-of-the-art test-time adaptation methods [8, 3, 7] for classification task usually discard statistics calculated during source training and estimate data distribution based on each batch of data separately. However, this way of calculating the statistics is flawed since the sample size from data is usually too small to correctly estimate the data distribution, especially for lower batch sizes.

Moreover, as presented in Figure 2, the magnitude of the distribution shift might vary between sequences. For some domain shifts, keeping the BN statistics of source data could be more beneficial. Therefore, there is a need for a method that adjusts statistics used in BN accordingly.

We adapt a part of the method developed by us [6] to semantic segmentation task and use BN statistics from source data to estimate BN statistics $\phi_t = (\mu_t, \sigma_t)$ at time step t during test-time by linearly interpolating between saved statistics from source data ϕ^S and calculated values from current batch ϕ_t^T :

$$\phi_t = (1 - \beta)\phi^S + \beta\phi_t^T \quad (2)$$

where β is a parameter that weights the influence of saved and currently calculated statistics.

We utilize the symmetric KL divergence as a measure of distance between distributions $D(\phi_{t-1}, \phi_t^T)$ to adjust the value of β accordingly to the severity of the distribution shift:

$$D(\phi_{t-1}, \phi_t^T) = \frac{1}{C} \sum_{i=1}^C KL(\phi_{t-1,i} || \phi_{t,i}^T) + KL(\phi_{t,i}^T || \phi_{t-1,i}) \quad (3)$$

where C is the number of channels. β_t at time step t is calculated as follows:

$$\beta_t = 1 - e^{-\gamma D(\phi_{t-1}, \phi_t^T)} \quad (4)$$

where γ is a scale hyperparameter.

To provide more stability for the adaptation, we take into account previous $\beta_{1:t-1}$ values and use an exponential moving average for β_t update:

$$\beta = (1 - \alpha)\beta_{t-1} + \alpha\beta_t \quad (5)$$

where α is a hyperparameter.

4.2. Entropy-based Pixel Filtering

Training feedback from entropy minimization might be noisy and unreliable, considering that the model predictions can be incorrect. Some of the state-of-the-art TTA methods [3, 4] used for the classification task filter images on which they adapt based on the entropy of the model’s predictions to make adaptation more reliable.

We decided to use a similar approach for the task of segmentation. However, discarding the whole images on which we update the model could be sub-optimal, considering the low number of images in sequences and segmentation task. Instead, we mask out single pixels in the process of calculating the loss based on the entropy of prediction for those pixels. Pixels with predictions having an entropy higher than the predefined, constant threshold are masked and do not participate in the backpropagation process. This way, we are able to adapt more robustly, disregarding uncertain predictions.

5. Experiments

5.1. Evaluation Metrics

The main metrics defined by the organizers are as follows:

$$overall = mIoU - 2 \times mIoU_{drop} \quad (6)$$

where $mIoU_{drop}$ is calculated as:

$$mIoU_{drop} = mIoU_{source} - mIoU_{target} \quad (7)$$

The $mIoU$ is based on combined predictions from all frames, $mIoU_{source}$ is a $mIoU$ from the first 20 frames of each sequence, and $mIoU_{target}$ is a $mIoU$ from 180th to 220th frame of each sequence. Apart from *overall*, we utilize a simple $mIoU$ metric in our experiments.

5.2. Implementation Details

We use the code repository provided by the challenge organizers for the development and evaluation of our method. Additionally, we implemented the *overall* metric locally ourselves, as it was only available on the evaluation server. The presented results are from our local evaluation on six sequences from the validation split unless stated otherwise.

We utilize DeepLabv3+ [1] with ResNet50 [2] backbone as a source model, with weights provided by the organizers. During TTA, we use a learning rate equal to 0.00006/4. The γ parameter value from Equation 4 is set to 0.1 and α from Equation 5 to 0.005, unless stated otherwise. The entropy threshold for discarding the pixels with uncertain predictions from the adaptation process is equal to $0.3 \times \ln 14$, where 14 is the number of classes and the $\ln 14$ represents the maximum entropy value.

5.3. Results

Table 2 presents the performance with different combinations of components of our method. It can be seen that adding each element to our base method (entropy minimization) increases both $mIoU$ and overall metrics. The most significant improvement is achieved in terms of overall metric by adding Dynamic BN Statistics Update to **B** configuration.

Table 2. $mIoU$ and *overall* metrics of different combinations of components of our method. The results are from the local evaluation on the validation split.

Method	$mIoU$	<i>overall</i>
A: TENT [7] baseline (entropy minimization)	56.9	39.0
B: A + Adapting Only Backbone’s BN Weights	57.7	41.3
C: B + Dynamic BN Statistics Update	58.6	49.1
Ours: C + Pixel Filtering	58.8	50.2

We show our final results from the server evaluation on test split in Table 3.

Table 3. Evaluation of the performance of our method from the evaluation server on the test split.

$mIoU$	$mIoU_{drop}$	$mIoU_{source}$	$mIoU_{target}$	<i>overall</i>
71.4	23.3	76.5	53.2	24.7

5.4. Ablation Study

Table 4 displays the performance of the baseline TENT [7] method, while different parts of a model are updated during test-time. The results show that adapting only BN weights of the backbone achieves the best results in terms of both metrics. Moreover, keeping all the weights of an adapted part unfrozen, instead of only BN ones, significantly degrades the plain entropy minimization performance.

Table 4. $mIoU$ and *overall* metrics of TENT [7] baseline with different configurations of adapted weights. The results are from the local evaluation on the validation split.

Adapted Weights			$mIoU$	<i>overall</i>
Backbone	Head	BN weights only		
✓			4.6	-101.4
	✓		18.7	-61.8
✓	✓		4.4	-88.3
✓		✓	57.7	41.3
	✓	✓	54.4	16.3
✓	✓	✓	56.9	39.0

Additionally, we explored different thresholds for filtering the unreliable pixel predictions for adaptation. Results are displayed in Table 5.

Lastly, in Table 6, we show the performance of our method with different parameters of dynamic BN statistics update, namely γ from Equation 4 and α from Equation 5.

Table 5. Performance of our method with different values of entropy threshold for filtering the unreliable pixel predictions. The results are from the local evaluation on the validation split.

Entropy threshold	$mIoU$	<i>overall</i>
$0.40 \times \ln 14$	58.6	49.2
$0.35 \times \ln 14$	58.6	48.9
$0.30 \times \ln 14$	58.8	50.2
$0.25 \times \ln 14$	58.8	50.0
$0.20 \times \ln 14$	58.8	49.4
$0.15 \times \ln 14$	58.7	47.5

Table 6. Performance of our method with different values of dynamic BN statistics update parameters - γ and α . The results are from the local evaluation on the validation split.

γ	α	$mIoU$	<i>overall</i>
1	0.5	58.0	40.3
	0.05	58.3	42.0
	0.005	58.3	43.8
	0.0005	58.6	46.0
0.1	0.5	58.2	45.0
	0.05	58.7	48.2
	0.005	58.8	50.2
	0.0005	58.8	47.5
0.01	0.5	58.2	43.8
	0.05	58.4	45.6
	0.005	58.6	49.6
	0.0005	58.8	47.9
0.001	0.5	58.1	43.3
	0.05	58.3	45.2
	0.005	58.8	48.0
	0.0005	58.8	48.0

5.5. Things That Didn’t Work

Apart from the techniques used in our final method, we experimented with more approaches that didn’t work.

We tried to adapt the model during test-time while keeping the features of the original and augmented images consistent by adding additional term to the loss function.

Moreover, we explored preserving a buffer of a low number of previous predictions and averaging them to obtain more reliable pseudo-labels. Additionally, to account for different positions of objects in images in between the frames, we considered using optical flow to unify the predictions into a single time step.

6. Conclusions

In this work, we present a brief analysis of the problem of continuous test-time adaptation and demonstrate our methods for semantic segmentation task. We were able to build upon and outperform the baselines.

Acknowledgement. Bartłomiej Twardowski acknowledges the grant RYC2021-032765-I.

References

- [1] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [3] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yaofo Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient test-time model adaptation without forgetting. In *International Conference on Machine Learning (ICML)*, 2022.
- [4] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Zhiqian Wen, Yaofo Chen, Peilin Zhao, and Mingkui Tan. Towards stable test-time adaptation in dynamic wild world. In *International Conference on Learning Representations (ICLR)*, 2023.
- [5] Tao Sun, Mattia Segu, Janis Postels, Yuxuan Wang, Luc Van Gool, Bernt Schiele, Federico Tombari, and Fisher Yu. SHIFT: a synthetic driving dataset for continuous multi-task domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [6] Damian Sójka, Sebastian Cygert, Bartłomiej Twardowski, and Tomasz Trzciniński. Ar-tta: A simple method for real-world continual test-time adaptation, 2023.
- [7] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *9th International Conference on Learning Representations (ICLR)*, 2021.
- [8] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.